Volume 09 Issue 04 Published, November 9, 2005 ISSN 1535-864X DOI: 10.1535/itj.0904

Intel® Technology Journal

Electronic Package Technology Development

Power Delivery for High-Performance Microprocessors

More information, including current and past issues of Intel Technology Journal, can be found at:
http://developer.intel.com/technology/itj/index.htm

Power Delivery for High-Performance Microprocessors

Kemal Aygün, Technology and Manufacturing Group, Intel Corporation
Michael J. Hill, Technology and Manufacturing Group, Intel Corporation
Kimberly Eilert, Technology and Manufacturing Group, Intel Corporation
Kaladhar Radhakrishnan, Technology and Manufacturing Group, Intel Corporation
Alex Levin, Digital Enterprise Group, Intel Corporation

Index words: power delivery, decoupling capacitors, Fmax, impedance profile

ABSTRACT

This paper provides an overview of the trends, challenges, and solutions associated with delivering power to high-performance microprocessors. Due to the large power levels in today's microprocessors, it is not uncommon to design a Power Delivery Network (PDN) with a sub-milliohm impedance target. Apart from the obvious design challenges, there are measurement challenges associated with characterizing these low-impedance power delivery networks. By using a combination of active and passive measurement techniques, it is possible to successfully characterize the power delivery performance of the system. Since most of the pre-silicon power delivery design decisions are made based on modeling data, it is important to have accurate, fully calibrated simulation models. The simulation models used to analyze the PDN for the Intel® Pentium® 4 processor are shown to have good correlation with the measurement results. While the importance of the ability to measure and model the PDN cannot be understated, it is equally important to fully comprehend the impact of power delivery noise on the overall system performance. This allows the system designer to make the right tradeoffs in maximizing performance without exceeding the cost budget.

INTRODUCTION

The number of transistors in a microprocessor chip has been growing exponentially in accordance with Moore's Law. Microprocessor current levels have been increasing rapidly as transistors get smaller and faster. Due to the large current levels in today's microprocessors, it is imperative to have a low-impedance path from the power supply to the die. Failure to do so can result in excessive noise that can impact performance by limiting the maximum operable frequency [1].

VR Components & LGA Socket on MB

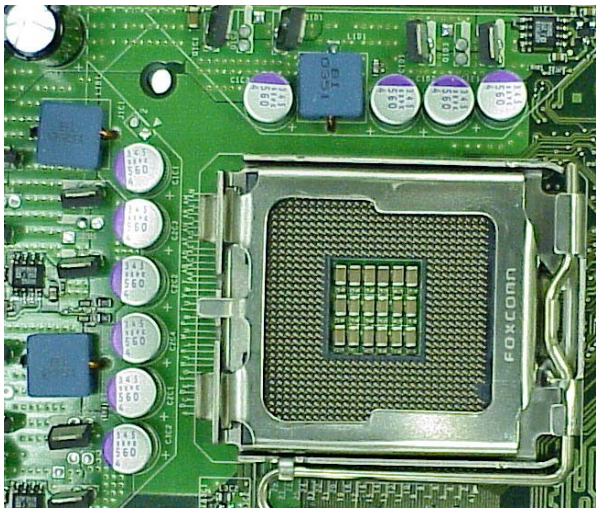

Package – Top View **Package – Bottom View**

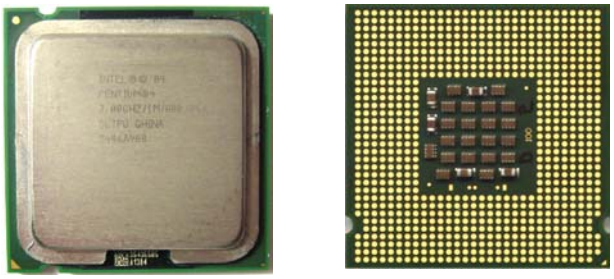

Figure 1: Power delivery solution for the Pentium® 4 processor

Since the current drawn by the processor can change suddenly, the impedance target needs to be met across a wide range of frequencies. This is typically accomplished by using a multi-stage decoupling solution with different types of capacitors. Figure 1 is a picture of the power

®Pentium is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

delivery solution for the Pentium 4 microprocessor on an LGA775 system. The high-frequency capacitors are placed on the land-side of the package. Ceramic capacitors are used for mid-frequency decoupling and are placed on the motherboard inside the socket cavity. Bulk capacitors are placed on the motherboard at the output of the voltage regulator to address the low frequency decoupling needs. These decoupling stages along with the voltage regulator and the on-die capacitance constitute the Power Delivery Network (PDN) of the microprocessor.

The primary objective of the power delivery designer is to pick the right type and location for these components in a way that will allow him or her to meet the impedance target. A majority of the power delivery design decisions such as the number of power and ground layers and the location of the capacitors is made prior to the availability of the first silicon. For this reason, these decisions have to be made based on data from the simulation models. Due to the sheer volume in the microprocessor market, the decision to add or remove a single capacitor can have a significant financial impact. As a result, it is important to have fully calibrated, reliable simulation models. Calibration of the simulation model is accomplished by collecting measurement data from the previous generation's microprocessor. Traditionally, power delivery validation was carried out in the time domain by monitoring the noise on the die sense lines while a high activity program is run on the processor. However, it is not convenient to use the time domain measurement data to calibrate the simulation models, due to the uncertainty associated with the die excitation model. In order to circumvent this issue, a frequency domain scheme is used to measure the impedance profile of the power delivery network as a function of frequency. This provides a metric that is independent of the die excitation and allows for direct correlation with the simulation model. In addition to the active measurements on a functional processor, passive characterizations on the individual components are often required to enhance the accuracy of the simulation models. For instance, the capacitance value specified by the vendor is often higher than the effective capacitance of the component under typical use conditions. A measurement scheme for measuring the effective capacitance as a function of temperature, DC bias, and AC signal level is described.

So far we have provided an overview of the power delivery problem for microprocessors. In the next section, we go over some of the power delivery trends that are seen in Intel microprocessors. We look at current and voltage trends as a function of time. We also cover some of the leakage power issues and discuss how they are driving the switch to multi-core processors. In section 3 we describe in detail the active and passive power delivery metrologies that are used to characterize the system power delivery performance. In section 4 we describe the construction of the simulation model and also include some model correlation results. Finally, in section 5 we focus on the impact of the power delivery noise on system performance.

POWER DELIVERY TRENDS

The number of transistors on a microprocessor chip has been increasing at an exponential rate. At the same time, these transistors have been switching faster to improve performance. These two trends combine to drive up the current consumed by microprocessors. Even though a part of this increase is offset by the reduction in the voltage levels and the transistor size, microprocessor current consumption has still been increasing at an exponential rate over the last two decades as shown in Figure 2. The brief respite in the current scaling in the mid-80s can be attributed to the switch from NMOS to CMOS technology.

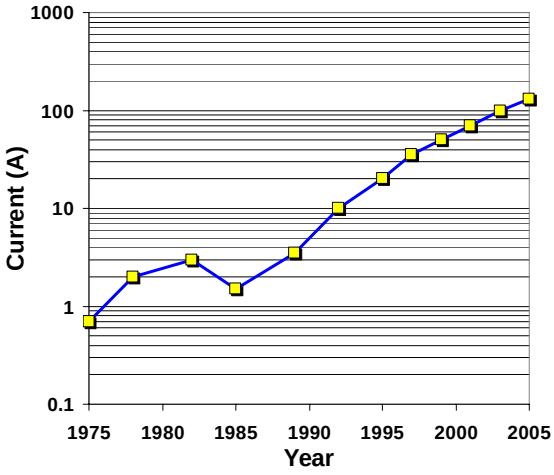

Figure 1: Microprocessor current trends in Intel microprocessors

As the dimensions on die get smaller to accommodate the increasing device density, the die voltage levels have been scaling down to meet oxide reliability constraints. Figure 3 shows the silicon feature size as a function of time. From the figure, we can see that the feature size has been scaling by a factor of ~0.7 x every two years. This corresponds to a doubling of the device density during the same period in accordance with Moore's Law. As the device dimensions continue to get smaller, the gate oxide thickness has gone from about 100nm back in the 1970s to about 1nm in today's process. In order to comply with the oxide reliability requirements, the die voltage has been scaling down as well as shown in Figure 3. The lowered operating voltage drives a lowered noise requirement. This trend coupled with the increasing current yields a

power delivery impedance target that is fast approaching sub-milliohm levels.

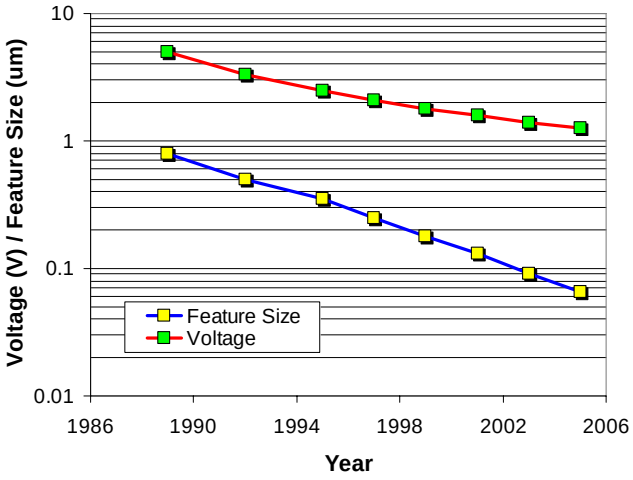

Figure 2: Microprocessor voltage and gate length trends

An unfortunate by-product of the reduction in the device dimensions is the increase in leakage power. Today's transistors start conducting current even when they are turned off and this current is referred to as leakage current. Figure 4 shows the growth in leakage power as a percentage of the total power supplied to the processor. From the figure, it is clear that the leakage power has grown from being negligibly small to being an appreciable percentage of the total power in a short period of time. If left unchecked, leakage power would soon exceed the active power consumption. One way to combat the leakage power issue is by slowing the frequency growth. With a reduced emphasis on the processor frequency, the process parameters can be tweaked to reduce leakage current at the expense of transistor switching speed.

With frequency no longer being the primary knob for improving the processor performance, system architects have turned to other avenues in an effort to improve the overall performance. One example of this is the switch to multiple cores. By adding an extra logic core and reducing the switching frequency, the processor can get a performance boost without a significant power penalty.

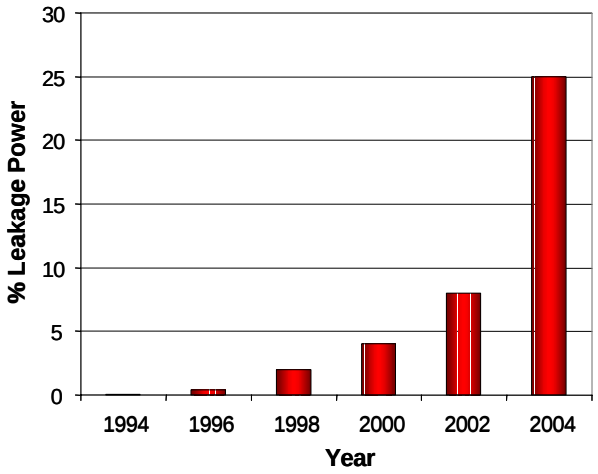

Figure 3: Leakage power growth

POWER DELIVERY METROLOGIES

The power delivery demands of today's microprocessors drive us towards bigger packages with more layers and better decoupling capacitors. The decision to add or remove a single component in a high-volume product can have a significant financial impact. As a result, it is important to have accurate methods to assess and quantify the impact of any changes made to the PDN.

Active Power Delivery Measurements

Traditionally active measurements for power delivery have often been performed in terms of voltage "droop" recordings [2]. In a droop measurement, the processor is periodically driven from a low-power consumption state to a high-power consumption state. This is also referred to as a large "*dI/dt*" event to emphasize the large instantaneous change in the amount of current drawn by the die. Then, the voltage at the die power rails, *v(t)*, is measured as a function of time during this transition. Given two different PDN solutions, the one with better performance will respond to the large *dI/dt* event better and this will be reflected in the amount of undershoot (i.e., droop) and overshoot in *v(t)*. The fundamental problem with this method is that it is usually very difficult to determine or measure the exact current drawn by the die, *i(t)*, during a typical droop measurement. Since *v(t)* is a strong function of *i(t)*, this essentially makes it considerably difficult to interpret and compare the droop waveforms obtained under different test conditions and for different processors.

A more direct assessment of the performance of the PDN can be achieved by measuring its impedance as a function of frequency. In order to measure the impedance profile of the PDN, we utilized the method described in [3], which is

also similar to the method presented in [4]. In this method, the clock tree of a microprocessor is directly pumped by an external clock signal. At the same time, the processor is held in a reset state to ensure that the processor's only toggling gates are those in the clock tree. This provides a current draw that can be directly controlled by the injected clock signal [3]. The procedure for measuring the transient die voltage and the transient die current for our particular implementation of this metrology has been detailed in [5]. Once *v(t)* and *i(t)* are measured, the impedance of the PDN as seen by the die can be computed as

$$Z(f) = V(f)/I(f), \qquad (1)$$

where $V(f)$ and $I(f)$ denote the Fourier transforms of *v(t)* and *i(t)*, respectively. For the stimulus used in this measurement it can be shown that *i(t)* is square pulse train, the frequency of which is equal to the frequency of the envelope of the injected clock signal [3]. Let this frequency be denoted by f_T. Then $I(f)$ can be represented as

$$I(f) = \sum_{k=-\infty}^{\infty} a_k \delta(f - kf_T) = \sum_{k=-\infty}^{\infty} I_k(f), \qquad (2)$$

where $\delta(\cdot)$ is the Dirac delta function, a_0 is the DC value of *i(t)*, $a_k = 0$ for even *k*, $a_k = -jI_{pp}/(k\pi)$ for odd *k*, and I_{pp} is the peak-to-peak magnitude of *i(t)*. Consequently, $V(f)$ can be expressed as

$$V(f) = \sum_{k=-\infty}^{\infty} Z(kf_T) I_k(f) = \sum_{k=-\infty}^{\infty} V_k(f). \qquad (3)$$

The impedance value at frequency kf_T can then be computed as

$$Z(kf_T) = V_k(f)/I_k(f). \qquad (4)$$

Based on equations 1-4, the impedance "extraction" algorithm at a given f_T can be summarized as:

1. Measure (or in this case compute [5]) $I_k(f)$ for $k = 1, \ldots, n$.
2. Measure $V_k(f)$ for $k = 1, \ldots, n$.
3. Compute $Z_k(kf_T)$ for $k = 1, \ldots, n$.

Here, *n* is the actual number of harmonics used in the measurement with the assumption that there is a measurable amount of energy in $V_n(f)$. Once $Z_k(f_T)$ for $k = 1, \ldots, n$ is extracted, the value of f_T can be changed and the procedure can be repeated at this new frequency. By varying the value of f_T, $Z_k(f_T)$ can be extracted over a broad range of frequencies. Note that based on the choice of f_T values and *n*, some of the kf_T values for different f_T will be the same. This forms a self-consistency check for the measurement. If the PDN to be characterized is linear, then the impedance values extracted using these different harmonics will be identical.

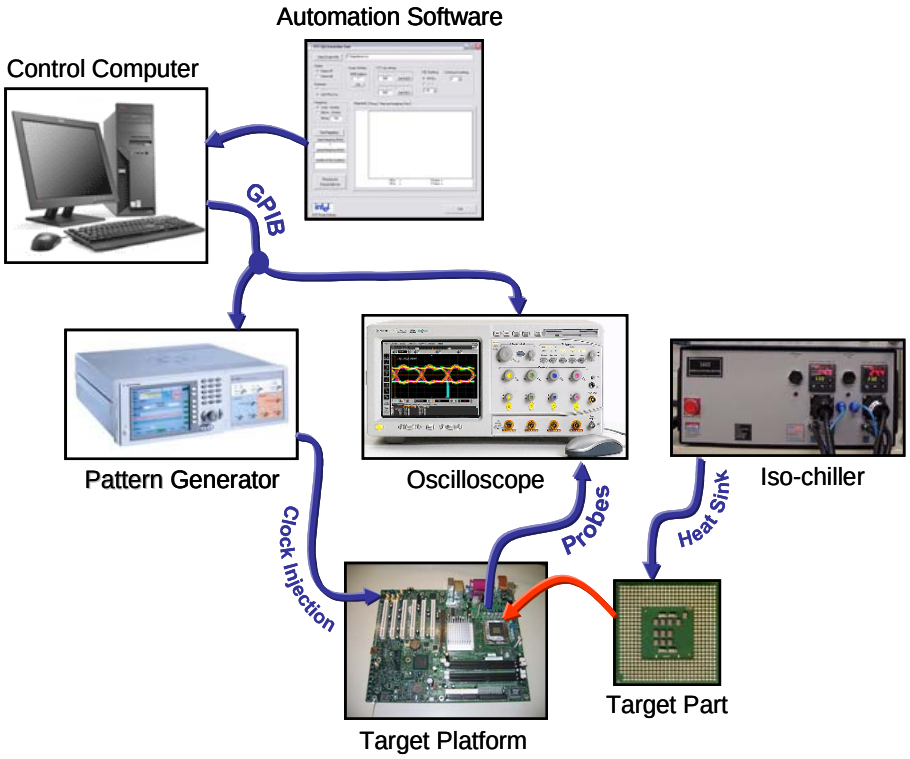

Figure 4: Equipment setup for impedance measurement

The equipment set up for the impedance measurement is illustrated in Figure 5. The injected clock signal is provided via a pattern generator. A high-bandwidth high-speed oscilloscope is used to measure *v(t)* and *i(t)*. The results in this paper were obtained using an Agilent AG31104* pattern generator and an Agilent AG54855A* oscilloscope with 1134A high-impedance probes [6]. Both the magnitude and phase of $V_k(f)$ were measured using the fast Fourier transform function of the oscilloscope. The measurements were performed under temperature controlled conditions using an USTC* iso-chiller.

There are two main challenges in measuring the impedance as described above. The first one is to devise a method to cleanly inject a high-frequency clock signal into the processor clock tree from an external pattern generator. The clock frequencies utilized in this measurement are on the order of a GHz and injecting them through motherboard and package into the silicon requires special care. A special test motherboard was designed for this purpose where the traces for clock injection were routed using the shortest possible path through the motherboard. Similar care was applied to the design of the sense lines where the *v(t)* (or $V_k(f)$) values were measured. The second challenge involves the labor associated with the measurement. Setting up the pattern generator and the oscilloscope manually at each frequency of the measurement to generate and measure the necessary waveforms is a very time-consuming task. To resolve this issue an automation tool was developed using Visual Basic*. At each frequency of the measurement the automation software communicates with the pattern generator and the oscilloscope through the General-Purpose-Interface-Bus (GPIB) and performs the necessary operations to compute the impedance. As a result of this automation the impedance profile of a typical PDN can be measured from Hz to hundreds of MHz in a matter of minutes. The impedance measurement technique was used to characterize the performance of various advanced decoupling solutions such as array capacitors [7]. One particular result from [7], which compares the performance of a package with an array capacitor to that of a package with standard package capacitors, is shown in Figure 6. As demonstrated by this figure the results obtained by this method unveil information regarding all the decoupling stages in the PDN in a very transparent way. This makes it very easy to compare different technologies in terms of their impact on the overall power delivery performance.

* Other brands and names are the property of their respective owners.

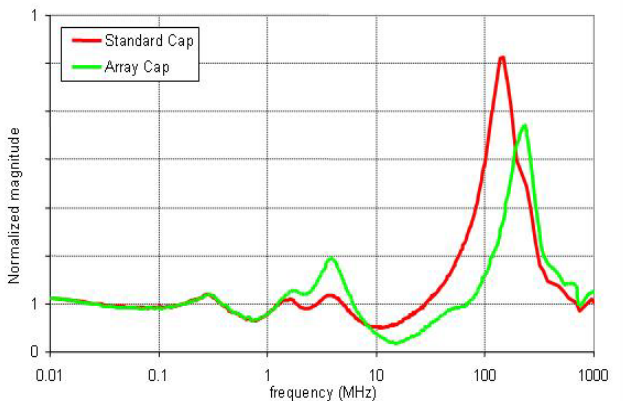

Figure 5: Measured impedance profiles for two packages with different capacitor solutions

Passive Power Delivery Measurements

Although characterization of the impedance profile of a functional microprocessor PDN provides a good picture of the performance of the complete PDN system, it is often desirable or necessary to examine individual components of a PDN outside of a functional system. Typically these passive measurements are helpful when there is a miscorrelation between measurement and modeling of the complete PDN, or when one desires to examine a new PDN component for use in the design of a system that is not yet available for active testing. In many of these cases the PDN component of interest is a high-performance decoupling capacitor. Typically these capacitors have ultra-low Equivalent Series Inductance (ESL) and Resistance (ESR), while maintaining high levels of capacitance. The space allotted to the PDN requires that the components have small form factors, and to achieve the desired high capacitance values, the internal spacing of the capacitor plates continues to decrease. Meanwhile the permittivity of the capacitor materials is also being driven higher. These two factors result in many of today's common PDN capacitors exhibiting non-linear behavior. Although non-linear capacitor behavior can be very complicated, to a first order approximation, the effective capacitance of a capacitor can be found by linearizing the capacitance around an operating point. Specifically, by measuring the capacitor's performance at different temperatures, DC biases, and AC signal levels, an effective value can be determined for a given use condition. This value can then be used in place of the manufacturer's specified value for modeling and for comparison of different capacitor solutions. This "use condition" value thus provides results that more closely resemble the real system performance because in a typical microprocessor system, the actual use condition is very much different from the conditions the capacitor experiences when the

industry-standard measurement techniques are used. Generally these industry-standard techniques involve capacitor measurement at room temperature with no DC bias, and they utilize an AC test signal that may be 1 volt or more. This is in stark contrast to the typical conditions seen in a CPU PDN. Most of today's CPU decoupling capacitors will experience temperatures that may exceed 50°C, with 1 to 2 V DC bias, and will never see AC signals above 100mV or so.

From a practical viewpoint, it is difficult to decide on one correct use condition that is appropriate for all microprocessors. Thus, simply changing the industry-standard measurement to a different set of conditions does not universally solve this problem. Additionally, in many situations multiple products are supported by a given package and decoupling technology generation. In these cases there may be many different use conditions for a given type of decoupling capacitor.

To enable practical use condition measurements an automated system is used to collect capacitor performance data at a large number of varying use conditions. In a typical case the system may examine the effective capacitance provided by a device as the temperature, DC bias, and AC signal levels are swept from 0°C – 100°C, 0VDC – 3VDC, and 5mVrms – 100mVrms, respectively. These data are then least squares fit to an 11-term, second-order polynomial in three variables (T, VDC, VAC), allowing for a compact representation of the large data set. This method was selected after studying the measured data for a number of different capacitor types, and it has been shown to be able to reproduce the original measured data with reasonable accuracy. End users of this information can simply load the coefficients describing the capacitor's performance into a custom calculator tool, or can manually calculate a result, to determine an appropriate effective capacitance for their application.

One system that utilizes this technique to characterize package and motherboard decoupling capacitors is shown in Figure 7. This system utilizes a GPIB instrumentation bus to link a control computer to an Agilent 4294A* impedance analyzer, a Trio-Tech TC1000* thermal control chuck, and a Stanford Research SR630* thermocouple reader. A custom software package has been developed to allow convenient adjustment of the sweep parameters and to perform the least squares fitting of the measured data.

* Other brands and names are the property of their respective owners.

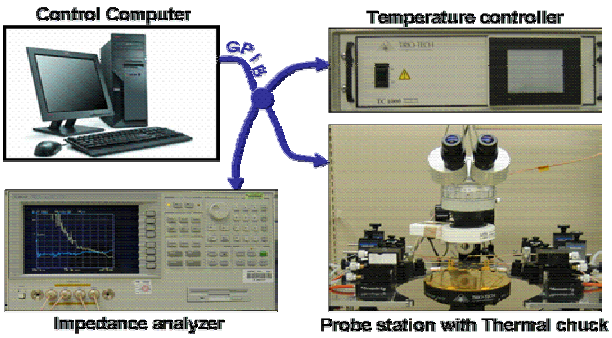

Figure 6: System for characterizing use condition capacitance

Figure 8 shows an example of a measurement performed on a sample capacitor. In this case the manufacturer, using the industry-standard measurement technique, has specified this capacitor as a 100uF device. The data shown in Figure 8 demonstrate that this capacitor achieves an effective capacitance of 100uF only at high AC test signal levels, and at moderate temperatures with no bias. At conditions more relevant to a microprocessor PDN, the effective capacitance is in the 80uF range or less.

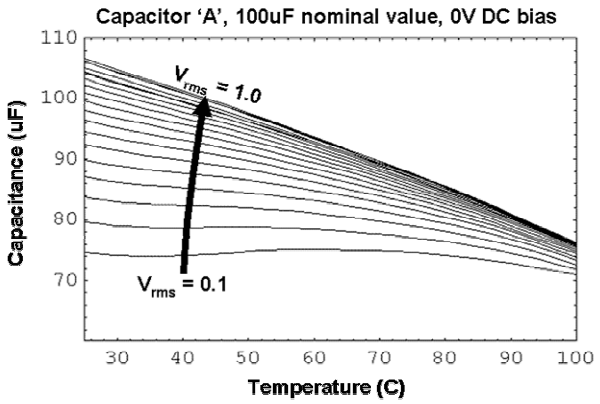

Figure 7: Effective capacitance of a 100uF capacitor as a function of temperature and AC signal level

With the significant variation in effective capacitance due to use condition parameters, it is important that PDN modeling activities utilize the use condition effective capacitance values; otherwise, the correlation between modeled and measured results will be poor.

MODELING THE POWER DELIVERY NETWORK

Since most design decisions are made prior to the availability of first silicon, and therefore a functional test system, it is important to have a good power delivery simulation model to assist with these decisions. In the past, simple lumped element models have been used to model the power delivery network. However, more

recently with the shrinking power delivery impedance targets, it becomes more important to include additional details that capture the non-uniform loading of the die and the spatial location of the power delivery components. For this reason, detailed distributed models of the package and the motherboard have replaced the traditional lumped element model for power delivery analysis.

This type of system simulation model is constructed using distributed circuit elements to represent the location of the power and ground planes in the package and the motherboard [8]. The model is created directly from the layout files of the package and the motherboard. The package portion is partitioned vertically into three sections along the z-direction and then split laterally into small cells. The required cell size is determined by the degree of resolution needed to resolve the particular package design. The top and bottom portions of this model represent the front-side and back-side build-up layers, respectively. The physical package layout in these sections is modeled using a quasi-static solver, Q3D*, from Ansoft [9]. The middle section of the model represents the core of the package. This area, which comprises the majority of the package height, is filled with large Plated Thru Hole (PTH) vias that connect the front-side planes to the back-side planes. These PTH vias are modeled as RL elements and their parasitics are also obtained from the quasi-static solver. Capacitors are treated as distributed RLC elements and attached to the package base or front-side depending on their location in the design. Values for the Equivalent Series Resistance (ESR) and capacitance of these components are determined from use condition passive measurements as described previously. The ESL of the capacitor is estimated using validated models. A pictorial view of the distributed power delivery model is shown in Figure 9. Based on the particular package topology, the resulting distributed model can include from thousands to tens of thousands of circuit elements.

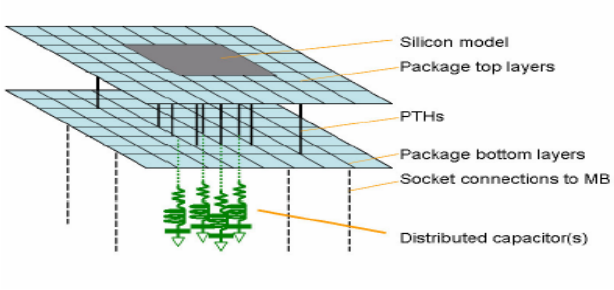

Figure 8: Pictorial view of the distributed power delivery model

When this model is used for time domain simulations, a complex Voltage Regulator (VR) model is attached to the motherboard and each die node is configured as a dynamic current load. For a frequency domain analysis, the VR is replaced by a power supply with an RL element in series. The R represents the static loadline of the VR and the L is proportional to the bandwidth of the VR. An AC voltage source is attached in place of the silicon stimulus at a single node, called the sense node. The model is then simulated across a wide range of frequencies using an AC sweep. The resulting voltage, which is proportional to impedance, is monitored at all the die nodes. The total system impedance, as seen at the sense node, can then be computed using the equation (5).

$$Z_{sense}(f) = \sum_{i=1}^{N_{die}} w_i \cdot Z_{i,sense}(f), \qquad (5)$$

In this equation, w_i represents the percentage of the total current that is drawn from the node *i*, and $Z_{i,sense}(f)$ represents the transfer impedance between node *i* and node *sense*.

The system impedance is computed from a weighted summation of the transfer impedances between the sense node and the remaining current sources. The weights, w_i, are determined by the current consumption of each silicon gird area *i*. To demonstrate the predictive capability of the method, a package capacitor de-population experiment was performed on a Pentium 4 processor [5]. The capacitors on the back-side of the flip chip pin grid array package were removed in the order shown in Figure 10. The corresponding impedance profiles were measured and simulated using the technique described above. Both the simulated and measured results are shown in Figure 11. The shifts in peak frequencies and the relative changes in peak height are in agreement for both the measured and simulated data. This demonstrates that the modeling approach is a robust and realistic characterization of the physical system.

* Other brands and names are the property of their respective owners.

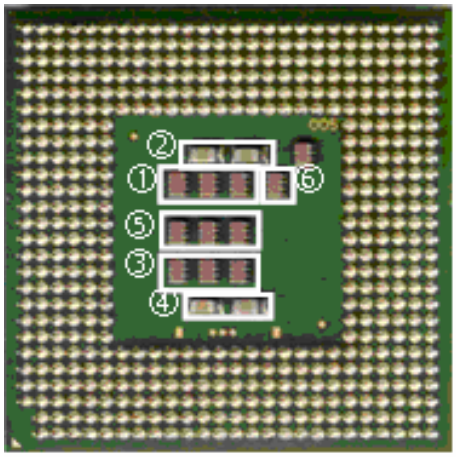

Figure 9: Ordering of capacitor removal

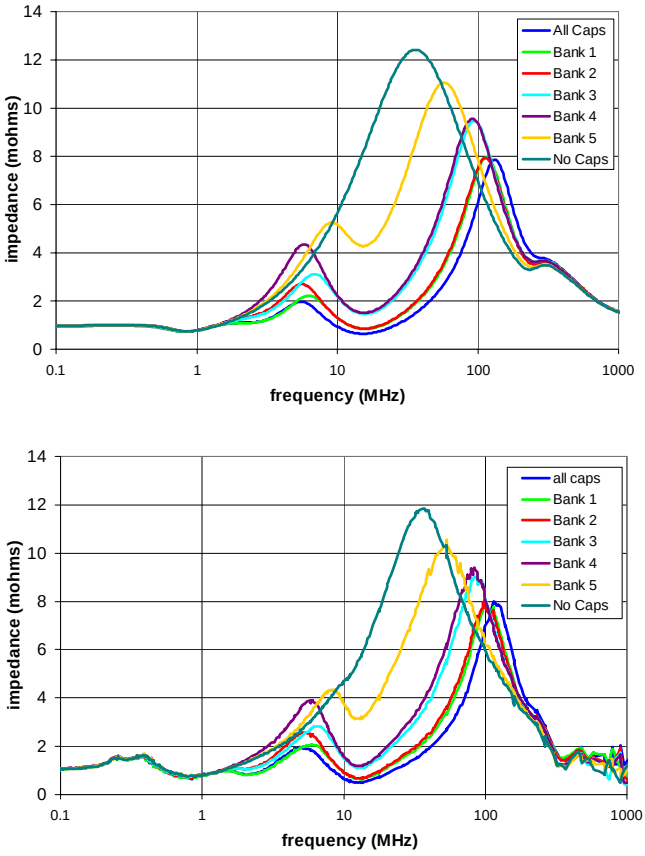

Figure 10: Simulated (top figure) and measured (bottom figure) results for the capacitor de-population experiment

Impact of PDN on System/Silicon Performance

These advances in measurement and characterization techniques have provided a significantly improved insight into power delivery network design and comparative performance. However, optimization of the network topology and component selection depends on an understanding of silicon performance in a less than ideal environment. Ultimately, one would strive to achieve purely resistive performance (flat Z-profile) for a power delivery network [10]. Regardless of the spectral content of the die stimulus, a fully flat network response will provide a constant relationship between current drawn by the silicon and voltage at the transistor level.

Silicon performance is typically defined in terms of logic path delay relative to a minimum clocking period. The delay is attributable to two physical contributors: transistor gate delay and RC wire delay. Transistor gate delay refers to the maximum switching time of cascaded devices, while RC wire delay describes the on-die interconnect contribution. The performance-limiting gate delay varies in direct proportion to on-die voltage at the location of critical circuits, and it is sensitive to PDN quality. The RC delay is temperature dependent and degrades under higher current/temperature conditions. Maximum operational frequency (Fmax) varies in direct proportion to the sustained voltage supply level at gate-limited timing paths, and inversely with temperature dependent RC delay [10]. Figure 12 shows the relationship between the change in Fmax and supply voltage (VID). At lower voltage levels, the gate delay tends to dominate and the relationship between Fmax and supply voltage is almost linear. However, at higher voltage levels, the RC delay begins to dominate and the curve begins to level off.

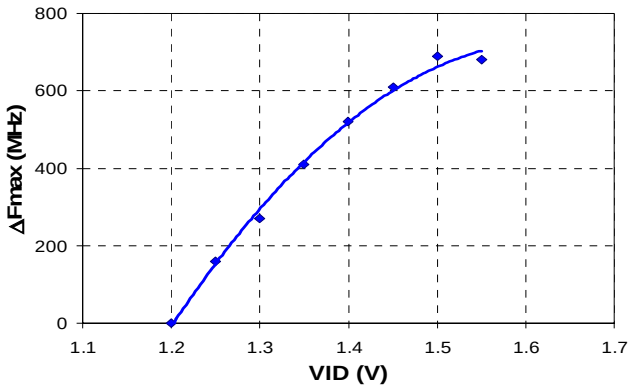

Figure 11: Plot of ΔFmax as a function of die voltage

If the voltage to current relationship through the PDN is linear (resistive) across all frequencies of network stimulus, the estimation of performance simply becomes an exercise in determining maximum transient current consumption. Worst-case current draw can be mapped to a voltage minimum via the network impedance. Critical circuit paths are then assessed for timing compliance (setup margin) at the minimum voltage, under temperature conditions associated with maximum current levels.

In a realistic, cost-effective PDN, the current/voltage/performance relationship is not so straightforward. Inadequate matching of board, package, and die capacitor characteristics promotes resonant behavior within the PDN. Cost constraints and manufacturability present significant obstacles to designing PDNs that are free of resonant behavior. At the resonant frequencies, current stimulus will excite a voltage response which is proportional to the associated impedance profile peak amplitude. For PDN resonances with periods much longer than the longest on-chip logic delays, the resonant noise effect on gate delay is akin to a DC voltage offset. To determine the performance impact from resonant network behavior, it is necessary to quantify:

a) the extent to which energy may be focused by modulating current at the frequencies of interest, without architectural restriction
b) the amplitude(s) of current which may be dynamically switched at the resonant frequencies

Understanding the component micro-architecture and the various power-savings modes inherent in it are necessary prerequisites. The next steps involve locating the highest current-consuming blocks within the component, and understanding the mechanisms for disabling those areas in order to create current "pulses." Global commands that can be issued to stop on-chip clock networks, or pause commands for high-speed circuit blocks, are obvious candidates. Once the highest current-consuming areas are located, and a disabling command is established, it is necessary to determine whether architectural considerations limit the frequency at which this disabling command may be asserted/deasserted. Intel microprocessors have been successfully instructed to draw a significant percentage of their maximum dynamic current at switching frequencies covering the region of resonant concern. By sweeping the frequency of stimulus to cover the network poles, CPU performance in the presence of the system's worst-case Vmin can be tested.

Flexibility in the commands issued to start and stop core activity provides several ways to determine the performance implications of peaks in the impedance profile. The simplest method is to target switching energy at a particular resonant frequency. Adding critical speed-path coverage into code that modulates core activity at various frequencies allows a measure of performance testing. By modulating core activity to induce worst-case voltage droops at the resonant frequency, the test can determine the relative maximum operating frequency (Fmax) for the part in that scenario. The Fmax result will not be absolute (since the testing is not exhaustive), but instead will indicate a relative reduction in speedpath performance, in comparison to code which executes coverage of the same speedpaths without significant focused current modulation.

The Peak Distortion Algorithm (PDA) is another method that is intended to provide alignment of the natural modes of the PDN to create a worst-case Vmin event [11]. A PDN impulse response (simulated or measured) and application of the PDA are used to determine the most damaging series of current-switching events that can be applied to the network. Through PDA, a stimulus pattern can be devised that aligns low, mid, and high frequency resonant behavior so that the voltage minima associated with each mode overlap in time. In a multi-pole power delivery network, the PDA-derived stimulus patterns are capable of creating lower voltage minima than with modulation focused at a single frequency. Figure 13 is an example simulation showing this effect.

Throughout these tests (whether focused on a single frequency or scripted to provide a worst-case pattern), commands are applied to limiting speedpaths within the core to check for timing failure conditions. A selected subset of known marginal test vectors is inserted into the high-activity loop in the executed code. Without modulation, the DC voltage setting of the platform can be lowered until failure of these test vectors is observed. The code is then run to stress the PDN (either targeted frequency or worst-case pattern). While running the test, the platform DC voltage can again be lowered until failure is observed. Assuming a worst-case current switching magnitude has been developed, the difference in DC setting between the non-modulated and modulated tests provides key insight into the voltage margin to failure that is lost because of network resonances.

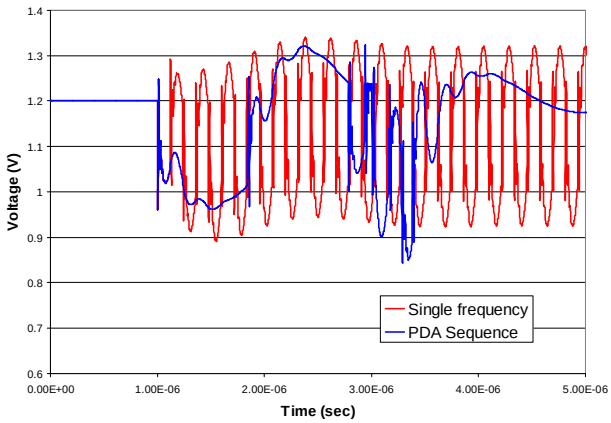

Figure 12: Voltage waveforms showing single-frequency switching at a network pole, and PDA-derived stimulus sequence inducing a lower voltage

Equipped with this relative performance information, the design engineer can pursue an informed course of PDN cost optimization. Capacitors can be depopulated or added, and values can be adjusted to suppress resonant behavior that has been demonstrated as damaging to overall performance. Through a regimen of PDN characterization and stress testing at the identified resonant frequencies, cost-effective decisions can be made that support core performance at minimal system expense.

CONCLUSION

Design and analysis of power delivery networks for microprocessors has become and will continue to be a challenging problem due to the trends in process and performance scaling. In this paper, we described some of the state-of-the-art measurement and modeling techniques as well as examined the impact of the PDN on the system performance. Measurement systems that can evaluate the performance of the PDN at both system and component level are necessary together with modeling techniques that have been validated against measured results. The success of an electrically efficient and cost-effective PDN analysis relies on the availability of these methods and an understanding of the impact of the PDN on the system performance.

REFERENCES

[1] A. Waizman and C. Chung, "Package capacitors impact on microprocessor maximum operating frequency," in the *Proceedings of the 51st Electronic Components and Technology Conference*, pp. 118-122, 2001.

[2] T. Rahal-Arabi, G. Taylor, M. Ma, J. Jones, C. Webb, "Design and validation of the core and IOs decoupling of the Pentium III and Pentium 4 processors," in the *Proceedings of IEEE: 11th Topical Meeting on Electrical Performance of Electronic Packaging*, pp. 249-252, 2002.

[3] A. Waizman, "CPU Power Supply Impedance Profile Measurement Using FFT and Clock Gating," in the *Proceedings of IEEE: 12th Topical Meeting on Electrical Performance of Electronic Packaging*, pp. 29-32, 2003.

[4] I. Kantorovich, C. Houghton, S. Root, J. St. Laurent, "Measurement of Low Impedance on Chip Power Supply Loop," *IEEE Transactions on Advanced Packaging*, vol. 27, pp. 10-14, Feb. 2004.

[5] K. Aygun, M. J. Hill, K. D. Eilert, and K. Radhakrishnan, "Measurement to Modeling Correlation of the Power Delivery Network Performance of a Microprocessor System," in the *Proceedings of IEEE: 13th Topical Meeting on Electrical Performance of Electronic Packaging*, pp. 221-225, 2004.

[6] www.agilent.com *

[7] K. D. Eilert, K. Radhakrishnan, M. J. Hill, and K. Aygun, "Use of Array Capacitors to Improve the Power Delivery Performance of a Microprocessor System," in the *Proceedings of 2005 ASME/Pacific Rim Technical Conference and Exhibition on Integration and Packaging of MEMS, NEMS, and Electronic Systems (InterPACK2005)*, 2005.

[8] K. Lee and A. Barber, "Modeling and Analysis of Multichip Module Power Supply Planes," *IEEE Transactions on CPMT*, part B, vol. 18, no. 4, pp. 628-639, Nov. 1995.

[9] www.ansoft.com *

[10] A. Waizman and C. Chung, "Resonant free power network design using extended adaptive voltage positioning (EAVP) methodology," *IEEE Transactions on Advanced Packaging*, vol. 24, no. 3, pp. 236-244, Aug. 2001.

[11] B. Casper, M. Haycock, and R. Mooney, "An Accurate and Efficient Analysis Method for Multi-GB/s Chip-to-Chip Signaling Schemes," in the *Proceedings of IEEE Symposium on VLSI Circuits*, pp. 54-55, 2002.

AUTHORS' BIOGRAPHIES

Kemal Aygün is a senior electrical packaging engineer at Intel Corporation in Chandler, Arizona. He has been with the Electrical Core Competency group working on power delivery analysis and metrology development since 2003. Kemal received his B.S. degree from Middle

East Technical University, Ankara, Turkey, and M.S. and Ph.D. degrees from the University of Illinois at Urbana-Champaign, all in Electrical Engineering. He is a member of the IEEE. Kemal can be reached by e-mail at kemal.aygun at intel.com.

Alex Levin is an analog engineer with Platform System Technologies at Intel Corporation in DuPont, Washington. He has been working on electrical design of Intel CPU and chipset packages since 2002. He joined Intel in 1997 where he has worked as an I/O circuit designer and as a signal integrity engineer. Alex holds a B.S. degree in Engineering Mechanics from the Johns Hopkins University and an M.S. degree in Electrical Engineering from the University of Washington. Alex can be reached by e-mail at alex.levin at intel.com.

Michael J. Hill is a staff electrical packaging engineer at Intel Corporation in Chandler, Arizona. He has been with the Electrical Analysis and Design group working on power delivery analysis and metrology development since 2002. Michael is a member of the IEEE Microwave Theory and Techniques Society, Etta Kappa Nu, and Tau Beta Pi. He holds B.S., M.S. and Ph.D. degrees in Electrical Engineering from the University of Arizona. Michael can be reached by e-mail at michael.j.hill at intel.com.

Kimberly Eilert is a packaging engineer at Intel Corporation in Chandler, Arizona. She has been with the Electrical Core Competency group working on power delivery analysis since 2003. She received her B.S. degree in Electrical Engineering from the California Institute of Technology in 2003. Kimberly can be reached by e-mail at kimberly.eilert at intel.com.

Kaladhar Radhakrishnan joined Intel in 2000 soon after receiving his Ph.D. degree in Electrical Engineering from the University of Illinois at Urbana-Champaign. He currently manages the Power Delivery Core Competency team within the Assembly Technology Development organization. Kaladhar can be reached by e-mail at kaladhar.radhakrishnan at intel.com.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

developer.intel.com/technology/itj/index.htm